

IN THE CLAIMS

Please amend the claims as follows:

1. (Previously Presented) A method of accessing shared memory in a computer system having a plurality of nodes, wherein each node includes a processor, a Translation Look-aside Buffer (TLB) associated with the processor and local memory wherein the local memory of each node includes a Remote Translation Table (RTT), the method comprising:

 distributing an application across the plurality of nodes building an application virtual address space in each of the plurality of nodes, wherein building an application virtual address space includes:

 building a local virtual address space for the application in each of the plurality of nodes, wherein the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node, wherein the virtual address generated by the application executing on the node includes a node number of the node;

 determining whether remote translation should be enabled; and

 if remote translation should be enabled, exporting the local virtual address space for each local node to the RTTs in each of the plurality of nodes, wherein exporting includes:

 requesting that the operating system on each node of the plurality of nodes enable remote translation; and

 shadowing the local virtual address spaces for each node across the plurality of nodes;

 if remote translation is enabled, translating a virtual memory reference to a physical memory address using the application virtual address space in the local RTT in each of the plurality of nodes, wherein translating the virtual memory reference includes

 translating in a source node the node number of the application virtual address into a node address associated with a remote node of the plurality of nodes and

 translating bits of the application virtual address into a physical page address for the

remote node; and

if remote translation is not enabled, sending the virtual memory reference to the remote node and translating the virtual memory reference into a physical address in local memory for the remote node using the local virtual address space in the RTT on the remote node.

2. (Previously Presented) The method of claim 1, wherein the local virtual address space is read from the Translation Look-aside Buffer (TLB) on the node.

3. (Original) The method of claim 1, wherein building an application virtual address space further includes performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes to complete exporting their respective local virtual address spaces.

4. (Previously Presented) A system comprising:

a plurality of nodes, each node including:

one or more processors having a Translation Look-aside Buffer (TLB);

a memory; and

a memory controller operatively coupled to the memory and the one or more processors, wherein the memory controller includes a Remote Translation Table (RTT), wherein the RTT is initialized upon the start of a process associated with an application by building virtual to physical address translations for local virtual address space in the node corresponding to the application, wherein each virtual address includes a node number of the node, and if remote translation should be enabled, by exporting the virtual to physical address translations for the local virtual address space from the node to the RTT in each of the plurality of nodes associated with that node through the application, wherein exporting includes:

requesting that the operating system on each node of the plurality of nodes enable remote translation; and

shadowing the local virtual address spaces for the node across the plurality of nodes;

wherein the TLB translates a virtual address reference received from the local

node into a physical address in the memory for the local node, if the node number for the virtual address is the local node number;

wherein, if remote translation is enabled, the RTT translates the virtual address received from the local node into a physical address in the memory for a remote node associated with the local node through the application using the virtual to physical address translations for the local virtual address space exported from the remote node; and wherein, if remote translation is not enabled, the RTT sends the virtual address received from the local node to the remote node and has the RTT on the remote node translate the virtual address into a physical addresses in the memory for the remote node using the virtual to physical address translation for the local virtual address space in the remote node.

5. (Original) The system of claim 4, wherein each of the plurality of nodes executes a synchronization operation that causes at least some of the plurality of nodes to wait for all of the plurality of nodes to complete exporting the virtual to physical address translations to their respective Remote Translation Tables.

6. (Previously Presented) A device-readable medium having instructions thereon that, when executed on a properly programmed information-processing device having a plurality of nodes, each node having one or more processors with a Translation Look-aside Buffer (TLB), a memory, and a memory controller coupled to the memory and the one or more processors, causes the information-processing device to perform a method comprising:

distributing an application across the plurality of nodes building an application virtual address space in each of the plurality of nodes, wherein building an application virtual address space includes:

building a local virtual address space for the application in each of the plurality of nodes, wherein the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node, wherein the virtual address generated by the application executing on the node includes a node number of the node;

determining whether remote translation should be enabled; and

if remote translation should be enabled, exporting the local virtual address space for each local node to the RTTs in each of the plurality of nodes, wherein exporting includes:

requesting that the operating system on each node of the plurality of nodes enable remote translation; and

shadowing the local virtual address spaces for each node across the plurality of nodes;

if remote translation is enabled, translating a virtual memory reference to a physical memory address using the application virtual address space in the local RTT in each of the plurality of nodes, wherein translating the virtual memory reference includes

translating in a source node the node number of the application virtual address into a node address associated with a remote node of the plurality of nodes and

translating bits of the application virtual address into a physical page address for the first remote node; and

if remote translation is not enabled, sending the virtual memory reference to the remote node and translating the virtual memory reference into a physical address in local memory for the remote node using the local virtual address space in the RTT on the remote node.

7. (Original) The device-readable medium of claim 6, wherein building a local virtual address space further includes performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes complete exporting their respective address space.

8. (Previously Presented) The device-readable medium of claim 6, wherein the local virtual address space is read from the Translation Look-aside Buffer (TLB) on the node.

9. (Previously Presented) A multinode system for implementing remote address translation, the system comprising:

a plurality of nodes, each of the plurality of nodes including:

one or more processors with a Translation Look-aside Buffer (TLB);

a memory; and

a memory controller operatively coupled to the memory and the one or more processors, wherein the memory controller includes:

a Remote Translation Table (RTT);

means for distributing an application across the plurality of nodes

means for building an application virtual address space in each of the plurality of nodes, wherein the means for building an application virtual address space includes:

means for building a local virtual address space for the application in each of the plurality of nodes, wherein the local virtual address space translates a virtual address generated by the application executing on that node to a physical address in local memory for that node, wherein the virtual address includes a node number of the node;

means for determining whether remote translation should be enabled; and

means for, if remote translation should be enabled, exporting the local virtual address space for each local node to the RTT in each of each the plurality of nodes, wherein means for exporting includes:

means for requesting that the operating system on each node of the plurality of nodes enable remote translation; and

means for shadowing the local virtual address spaces for each node across the plurality of nodes;

means for translating the virtual memory reference into a physical address in local memory using the TLB on the local node, if the node number is the local node number;

means for, if remote translation is enabled, translating a virtual memory reference to a physical memory address using the application virtual address space in the local RTT in each of the plurality of nodes, wherein means for translating the virtual memory reference includes

means for translating in a source node the node number of the application virtual address into a node address associated with a remote

node of the plurality of nodes and translating bits of the application virtual address into a physical page address for the remote node; and
means for, if remote translation is not enabled, sending the virtual memory reference to the remote node and translating the virtual memory reference into a physical address in local memory for the remote node using the local virtual address space in the RTT on the remote node.

10. (Previously Presented) The multinode system of claim 9, wherein means for building an application virtual address space further includes means for performing a synchronization operation that causes at least some of the plurality of nodes to wait for all nodes to complete exporting their respective local virtual address spaces.

11. (Currently Amended) A multi-node system for implementing remote address translation, the system comprising:

a network;
a source node coupled to the network, wherein the source node includes a first remote-translation table (RTT); and
a remote node coupled to the network, wherein the remote node includes a second RTT;
wherein on the remote node the second RTT is built using a first local address space on the source node exported from the source node to the remote node using an operating system call to perform the export;

wherein on the source node the first RTT is built using a second local address space on the remote node exported from the remote node to the source node using the operating system call to perform the export;

wherein the operating system enables remote translation utilizing the first and second RTTs, wherein enabling remote translation utilizing the first and second RTTs includes having the remote node translate a virtual memory address associated with the source node to a physical address on the source node as a function of the second RTT and having the source node translate a virtual memory address associated with the remote node to a physical address on the remote node as a function of the first RTT; and

wherein both the first and second RTTs include one or more virtual address addresses and each virtual address includes a node number of a remote node that built the virtual address.

12. (Previously Presented) The method of claim 1, wherein requesting the operating system on each node of the plurality of nodes enable remote translation passes control of the RTT to the operating system.

13. (Previously Presented) The method of claim 12, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.

14. (Previously Presented) The system of claim 4, wherein requesting the operating system on each node of the plurality of nodes enable remote translation passes control of the RTT to the operating system.

15. (Previously Presented) The system of claim 14, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.

16. (Previously Presented) The device-readable medium of claim 6, wherein requesting the operating system on each node of the plurality of nodes enable remote translation passes control of the RTT to the operating system.

17. (Previously Presented) The device-readable medium of claim 16, wherein passing control of the RTT to the operating system causes the operating system to maintain coherency of the RTT.

18. (Previously Presented) The method of claim 1, wherein requesting the operating system on each node of the plurality of nodes enable remote translation handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.

19. (Previously Presented) The system of claim 4, wherein, when remote translation is enabled, the operating system on each node of the plurality of nodes handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.

20. (Previously Presented) The device-readable medium of claim 6, wherein requesting the operating system on each node of the plurality of nodes enable remote translation handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.

21. (Previously Presented) The system of claim 9, wherein, when remote translation is enabled, the operating system on each node of the plurality of nodes handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.

22. (Previously Presented) The system of claim 11, wherein, when remote translation is enabled, the operating system on each node of the plurality of nodes handles requests to changes the application virtual address space configuration on a node-local basis, wherein handling requests includes disallowing an attempt to modify the application virtual address space outside scope of the local node.